

CLAIMS

What is claimed is:

- 5 1. A distributed system, comprising:
a set of nodes that communicate via a network;
a set of node applications distributed among the
nodes;
means for generating a time-stamp record for
10 each of a set of significant events associated with
one or more of the node applications such that the
time-stamp records provides a synchronized time base
across the nodes for the significant events.
- 15 2. The distributed system of claim 1, wherein the
means for generating a time-stamp record in one or
more of the nodes include a synchronized clock.
- 20 3. The distributed system of claim 2, wherein one
or more of the nodes include means for reading a time
value from the corresponding synchronized clock and
means for writing the time value into a local event
log that holds the corresponding time-stamp records.
- 25 4. The distributed system of claim 3, wherein one
or more of the nodes include ^{node dependent} means for generating an
event code for each significant events associated
with the corresponding node applications.
- 30 5. The distributed system of claim 4, wherein one
or more of the nodes include means for writing the

~~event code into the local event log along with the
time value.~~

42
5 6. The distributed system of claim 1, wherein the
means for generating a time-stamp record in one or
more of the nodes include a companion node having a
synchronized clock.

10 7. The distributed system of claim 6, wherein one
or more of the nodes include means for reading a time
value from the synchronized clock in the companion
node and means for writing the time value into an
event log that holds the corresponding time-stamp
records.

15 8. The distributed system of claim 1, further
comprising means for obtaining the time-stamp records
from the event logs via the network and analyzing the
time-stamp records using the synchronized time base.

20 9. The distributed system of claim 1, further
comprising means for starting and stopping the
generation of time-stamp records in one or more of
the nodes.

25 10. A method of performance monitoring in a
distributed system, comprising the steps of:
determining a set of significant events
associated with a distributed application in the
30 distributed system;

providing each of a set of nodes applications associated with the distributed application with the functionality to generate a time-stamp record when one of the significant events occur;

running an experiment in the distributed application that generates one or more of the significant events;

obtaining the time-stamp records from the node applications and analyzing the time-stamp records.

10

11. The method of claim 10, wherein each time-stamp record includes an event code associated with the corresponding significant events.

15

12. The method of claim 10, wherein the step of analyzing the time-stamp records comprises the step of generating a graphical representation of the time-stamp records.

20

13. The method of claim 10, further comprising the step of determining a set of delays in execution of the node applications associated with the generation of the time-stamp records.

25

14. The method of claim 13, further comprising the step of correcting the time-stamp records in response to the delays.